

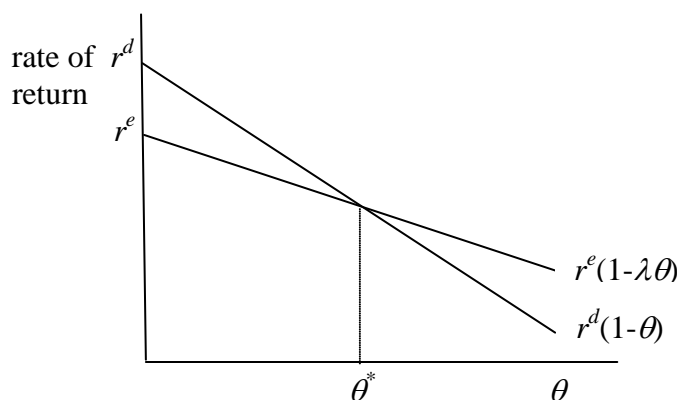
Lecture Note 10: Saving and Portfolio Responses to Taxation:
Introduction; Retirement Saving

Until now, our discussions of capital income taxation have treated saving as a single activity and assumed that any taxes imposed on capital income are applied as the income accrues. In reality, many types of assets serve as vehicles for saving, and the taxes imposed on saving often vary by asset (e.g., debt vs. equity vs. housing) or purpose (e.g., short-term saving vs. retirement saving). Also, the method of imposing taxes on capital accumulation varies, including taxation on accrual of income but also taxation on sale of assets (in the case of capital gains) or death (in the case of estate and inheritance taxes). Heterogeneity of tax treatment means that there are potential behavioral responses to taxation other than simply in the amount saved, and that questions of incidence and efficiency are also more complicated. Observed responses to the taxation of saving often reveal patterns of behavior that deviate from simple models of rational choice, deviations that affect not only the responses to policy but also the design of optimal policy.

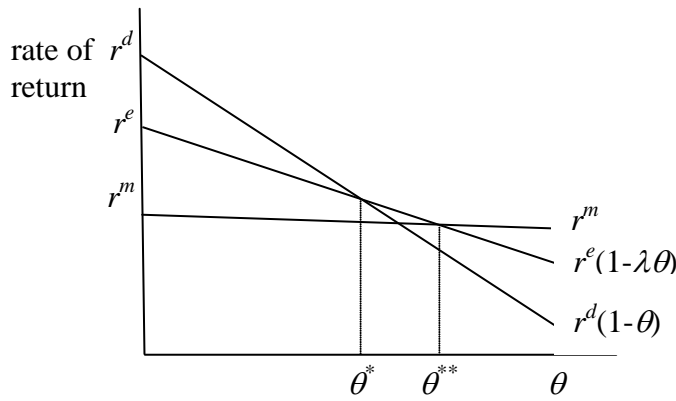
Taxation and Portfolio Choice

Under a progressive income tax, some individuals may have higher marginal tax rates than others on capital income. This difference in marginal tax rates may affect the level of individual saving, but will it affect the *composition* of saving, in terms of assets? There are two potential reasons why the composition of assets, i.e., the individual portfolio choice decision, may be influenced by taxation. The first is that not all assets are taxed at an individual's normal tax rate. The second, discussed below, is that taxes affect not only expected returns, but also the riskiness of returns. For a variety of reasons, some assets may be tax favored, that is, face a lower rate of tax than the individual's regular marginal tax rate. An important example is assets that deliver their income in the form of capital gains, which are taxed less heavily than other income in most countries. In the United States since 2003, lower tax rates also apply to dividends, so that all income from investment in equity, both dividends and capital gains, is tax favored relative to income from fully taxed assets, such as debt issued by companies or by the federal government.

To understand the impact of the existence of tax-favored assets, suppose there are two assets that are perfect substitutes from an investor's perspective except in the way they are taxed. Income from debt (i.e., interest) is taxed at the taxpayer's full rate of tax, θ , while income from equity is taxed at rate $\lambda\theta$, where $0 < \lambda < 1$. If r^d and r^e are the before-tax returns to debt and equity, then



the relative returns for investors are as shown in this graph. Investors with regular tax rates above θ^* will prefer equity, those with lower tax rates will prefer debt, and those at θ^* will be indifferent. Thus, we'd expect tax-favored assets to end up in the portfolios of high-bracket investors. If we added a third asset that is even more tax-favored, for example municipal debt, which is tax exempt, the picture would look like:



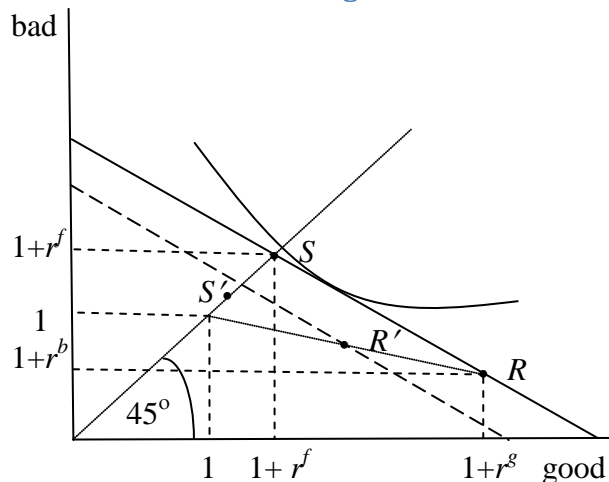
There will now be three groups of investors: those in the highest brackets, above θ^{**} , holding municipal debt; those between θ^* and θ^{**} holding equity; and those below θ^* holding taxable debt.

Note that we have drawn the graphs assuming that the more tax-favored the asset, the lower its before-tax rate of return. This must be true if positive amounts of all three assets are to be held. Otherwise, all investors would prefer tax-favored assets.

This translation of tax-favored status into lower before-tax rates of return is another illustration of tax capitalization, in this case the capitalization of tax benefits. For the marginal investor on the borderline between two assets, the tax benefits of the more lightly taxed asset are just offset by the asset's higher price (i.e., lower before-tax rate of return). How much before-tax rates of return differ depends on the relative asset supplies. For example, if the supply of municipal debt were reduced, then a higher value of θ^* , with fewer investors holding municipal debt, would clear the market for municipal debt. Put another way, in this model the tax benefits that are capitalized are those of the marginal investor, whose identity depends on asset supplies. Note that in this model the incidence of taxation comes in two components, through capitalization and direct taxation. For example, individuals who hold equity bear some tax through a lower before-tax rate of return than on debt, and some through their own (favorable) taxation of equity returns. Individuals who hold municipal debt bear taxes only indirectly, through capitalization. This distinction between capitalized taxes and taxes directly paid is important to keep in mind when looking at statistics on tax burdens that reflect only the latter.

So far, our characterization of tax-induced portfolio choice is unrealistic in predicting that each individual will specialize in a particular asset, which is clearly at odds with actual portfolio patterns. This is because assets typically differ in another dimension as well – their risk profiles. Differences in risk tolerance and a desire for portfolio diversification will also influence portfolio choice, and there will also be an interaction between taxation and risk, because taxes tend to dampen return fluctuations – after-tax returns have a lower variance than before-tax returns.

Taxation and Risk-Taking



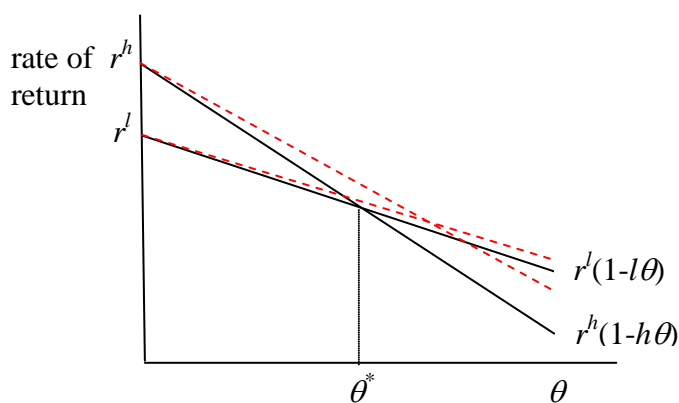
To consider the effects of taxation on investment in risky assets, consider a model in which there are two states of nature (“good” and “bad”) and two assets, one safe (S), with return r^f in both states, and one risky (R), with returns $r^g > r^f$ and $r^b < r^f$ in the good and bad states. (Note that risk aversion requires that in equilibrium the expected return on the risky asset exceeds that on the safe asset.) The two assets together define the budget line as shown, and the tangency illustrated in the figure corresponds to a portfolio with positive holdings of both assets.

If income taxes at a single rate are imposed on all returns, this will contract the points R and S toward the point $(1,1)$ on the 45° line and lead to a parallel shift in the budget line and the new points R' and S' as shown, with relative asset demands influenced by an income effect. However, taxes may also affect the slope of the budget line; assuming that the risky asset is equity and the safe asset is debt, lower individual taxes on equity income would favor the risky asset, but capitalization of the generally more favorable equity taxation into lower before-tax equity returns would favor debt. Which effect dominates depends on whether the individual is a low-bracket or a high-bracket taxpayer (i.e., has a value of θ below or above θ^* in the first figure above). As a consequence, a substitution effect will push higher-bracket taxpayers toward holding more of the risky asset, and will push lower-bracket taxpayers toward holding less, but diversified portfolios will still be in order. How much of each asset is held will also be affected by risk aversion. For example, an infinitely risk-averse investor would have Leontief indifference curves with a kink-point on the 45° axis, and hence would hold only the safe asset. A risk-neutral investor, on the other hand, would have straight-line indifference curves and choose to hold only the asset with the higher expected after-tax return.

Two other important points are worth making at this point about taxation and risk-taking, both of which can be related to the above two-state figure:

- (1) Tax systems typically treat gains and losses asymmetrically. While positive income is taxed, negative income (i.e., losses) does not receive a full tax refund. This means that the before-tax return on the risky asset in the good state may be $(1-\theta)r^g$, while the before-tax return in the bad state may simply be r^b , if $r^b < 0$, as is the case in the above figure. This would cause point R to shift horizontally to the left under taxation, to a point directly below R' , and hence to steepen the budget line and discourage investment in the risky asset.
- (2) Taxing risky assets reduces both expected returns and risk; the former discourages investment in risky assets, while the latter makes them more attractive. In one special case, the latter effect must dominate. Suppose that the tax system does not tax capital income generally, but just the excess over the safe rate of return, $(r^g - r^f)$ or $(r^b - r^f)$. Also assume that the tax system is symmetric, so that the issue just discussed does not arise. Then in the above two-state figure, taxation does not move point S , and simply shifts point R along the original budget line toward point S . For example, if the tax rate is 50%, R will move half-way from its original position to S . This does not change the investor's budget line, but must increase the portfolio share held in the risky asset. That is, a tax on *excess returns* – returns to risk-taking in excess of the safe rate of return – reduces a risky asset's expected return but in a way that does not change the investor's options and that encourages risk-taking. (A corollary is that the expected tax payment, which is positive because the expected before-tax return on the risky asset exceeds that on the safe asset, is of zero value to the investor and imposes no burden.) While this encourages private risk-taking, it also increases the risk borne by the government, unless the government can pool the risks of individual tax payments; the distribution of that risk by the government to individuals may in turn reduce investor risk-taking (since they will already have some risk in their lives) and might even undo the initial increase in risk-taking, a point made by Gordon (*QJE* 1985).

Evidence on the influence of taxes on portfolio choice is somewhat mixed but generally consistent with the theory that taxes should influence the mix of assets held; see Poterba (pp. 1126-1131). A recent application is in the paper by Kawano, which studies the impact of the 2003 reduction in the US rate of dividend taxation, mentioned above. The effect of the legislation was not only to make equity more attractive, relative to other assets, but also to increase the attraction of equities with high dividend yields (and hence a larger share of their income coming in the form of dividends, relative to capital gains) relative to equities with low dividend yields, especially for investors in higher tax brackets. Thus, we should have expected higher bracket investors, relative to low-bracket investors, to shift their portfolios toward high-dividend-yield stocks. In terms of the simple two-asset graph above, we can imagine the two assets being high-yield and low yield stocks, with the tax rates on both stocks depending on the individual's ordinary tax rate, θ , but with high-yield stocks facing a higher tax rate, since dividends faced a higher tax rate than capital gains prior to 2003. (Even since 2003, other provisions make the effective capital gains tax rate somewhat lower than the dividend tax rate, but the gap is much smaller than before.) The 2003 legislation lowers the tax rate on both assets, but it lowers the tax rate more for high-yield stocks, a relative benefit most valuable for those in high brackets.



As shown in the graph to the left, this change shifts the point of indifference to the right, given no change in before-tax returns. (We might also expect the before-tax return on high-yield stocks, r^h , to fall relative to that on low-yield stocks, r^l , in order to clear the market for the two types of stocks.) Indeed, Kawano finds a shift in portfolio sorting, with higher-bracket investors shifting more strongly toward high-yield stocks than low-bracket investors.

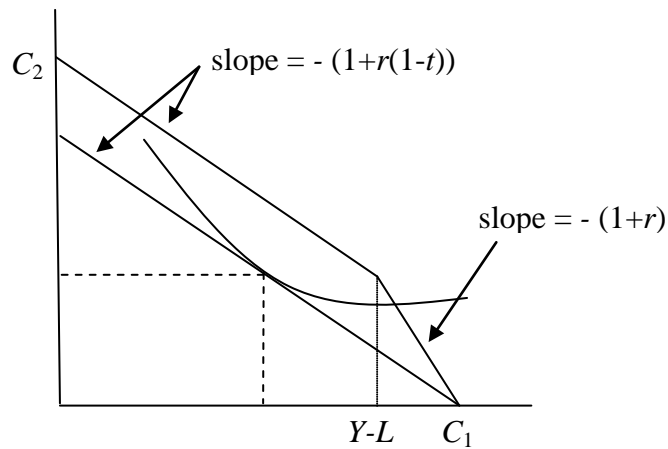
Tax-Favored Retirement Saving

As already discussed, one way of implementing a consumption tax is to allow a deduction for saving from an income tax. Although the United States does not treat savings in general this way, it does provide such treatment for specially designated retirement saving, through employer-based accounts (known as 401(k) plans) and individual accounts, known as Individual Retirement Accounts (IRAs). The same is true in many other countries. Although the traditional tax treatment of such accounts involves the “consumption tax” approach – deduction of contributions to accounts and taxation of withdrawals – an alternative approach (in the United States, called Roth IRAs and Roth 401(k)s) is to ignore contributions and withdrawals as well as not taxing accruing income – as capital income would be treated under a labor income tax. These alternative approaches are often summarized as “EET” and “TEE” methods of taxation (T for taxable, E for tax-exempt), where the first letter corresponds to the treatment of contributions, the second to accruing income, and the third to withdrawals. As we discussed earlier, these two approaches both eliminate the tax on new saving if tax rates are constant over time. (As we also discussed, under a comprehensive tax that applies to all saving, the two approaches differ as well

in the treatment of existing assets, if withdrawals of preexisting assets are taxed under the consumption-tax approach, but this distinction does not arise here, for the special tax treatment applies only to assets in the specified accounts, not to other assets.) The main reason for the introduction of Roth accounts in the United States was due to budget accounting – to defer the revenue consequences of the tax reductions to future years, which would have conveyed no political advantage had such consequences been fully taken into account.

Tax-favored retirement accounts also typically have annual caps on the level of contributions, which would constrain individuals who wish to contribute large amounts. Also, contributions do not have to represent new saving – individuals can contribute by transferring existing assets from regular taxable accounts to retirement saving accounts, either directly (in the case of IRAs) or indirectly (in the case of 401(k)s) by reducing balances in regular taxable accounts in order to finance current consumption while at the same time having contributions to 401(k) plans deducted from their salaries by employers.

Thus, tax-favored saving accounts differ from the tax treatment of savings under a consumption tax or a labor income tax in three important respects: (1) they apply only to saving for retirement, and have various other provisions (such as early withdrawal penalties) that make them less attractive for saving for pre-retirement consumption; (2) they have annual contribution limits; and (3) they apply to balances in specified accounts, rather than all assets. In light of these differences, what should we expect the impact on saving to be? Consider first a simple two-period life-cycle example, in which individuals work in the first period and save what is not initially consumed for retirement consumption. Without special retirement saving incentives, the individual faces the original budget line, with slope $-(1+r(1-t))$, where t is the capital income tax rate, and chooses the point of tangency shown.



When the accounts are introduced, the individual can save up to limit L in tax-favored form, receiving a rate of return of r ; further saving still faces capital income taxation. Thus, there will be only an income effect if desired saving exceeds L . The same logic applies in the case where individuals have pre-existing assets, in which case they will hit the contribution constraint as long as their existing wealth plus new saving exceeds L .

On the other hand, some factors may make it more likely that contributions to accounts have more than income effects. First, to the extent that withdrawal restrictions make retirement saving an imperfect substitute for other saving, individuals may not wish to transfer all assets into retirement saving accounts; thus, they may be on the steeper slope of the outer budget line in the figure. Also, behavior beyond the standard rational choice assumptions may lead individuals to save more if there is a specific reason for saving in these accounts or if the saving is coordinated by employers. Empirical evidence is somewhat mixed, in part because it is difficult to come up

with clean natural experiments to assess the treatment of saving incentives, because “treatment” has typically differed across individuals only because of differences in characteristics that may also be associated with unobserved differences in saving propensities. For example, individuals who work for employers offering 401(k)s may differ from those who work for employers not offering 401(k)s. See Bernheim’s *Handbook* chapter, section 4. Gelber’s paper is a recent attempt to control for unobserved differences among individuals, by considering individuals who had to wait to qualify for 401(k) plans. In general, the evidence does suggest that retirement saving responds positively to favorable tax treatment. There is also considerable evidence that individuals are influenced by employer education efforts (as discussed by Bernheim, section 5.4). Recent research also suggests that individuals are strongly influenced by “default” policies of employers – that they are much more likely to participate in retirement saving if that is the default employers present, rather than if they must make the decision to enroll. Another strategy for inducing employees to participate is to require an active decision whether or not they wish to participate rather than giving them the option of passive non-participation. Carroll et al. compare active decisions and defaults. While both strategies increase participation, the authors argue that requiring an active decision is preferable to a specific default if tastes for saving are heterogeneous, for these tastes will be reflected in people’s decisions; but this conclusion will not hold if financial illiteracy provides an argument for overriding individual preferences.

Another interesting decision is what investments to hold in tax-favored saving accounts. In particular, if individuals have both tax-favored accounts and taxable accounts, how should their portfolio allocations vary across the accounts? One’s first intuition might be that individuals should base each account’s composition on the account’s tax treatment and the individual’s risk tolerance. For example, an individual who is not too risk averse might choose to hold a lot of equity in a taxable account, and less equity in a tax-favored account, since (following the logic of the above discussion of portfolio choice), the tax-favored account does not derive any additional benefit from the favorable tax treatment that equity receives. But this logic misses an important distinction: in the present case, the same person is holding the two accounts, and should optimize simultaneously while choosing their composition. This leads to a simple arbitrage argument that individuals should concentrate lightly taxable assets in taxable accounts and more fully taxable assets in tax-favored accounts. The demonstration of this result proceeds by contradiction. Suppose that there are two assets, for example fully taxed debt and partially taxed equity, and that the investor has both a tax-favored account and a taxable account, each with some equity and some debt. Then the individual can increase debt in the tax-favored account while reducing debt in the taxable account, offsetting the changes in each account with equal changes in equity, in such a way that the overall portfolio risk is unchanged but the expected after-tax rate of return is increased. See Poterba’s *Handbook* chapter, section 4.3. This arbitrage argument does not hold exactly because of the withdrawal restrictions on tax-favored retirement accounts, but allowing explicitly for this factor does not have an important impact; see the paper by Dammon, Spatt and Zhang. Nevertheless, as Poterba discusses, evidence suggests that individuals adopt similar asset allocations in their tax-favored and taxable accounts. Such behavior might be due to a tendency of individuals to maintain different “mental accounts” for different types of saving, but this is an open question.